IBM WebSphere MQ Low Latency Messaging, Version 2.0

High throughput and low latency to maximize business responsiveness

IBM WebSphere® MQ Low Latency Messaging is a messaging transport that is highly optimized for the very high-volume, low-latency requirements of financial services firms. Applications include the high-speed delivery of market data, trade data, reference data and event data in or between front-, middle- and back-office operations. Characteristic applications require extremely low (submillisecond) latency and high message volumes (ranging from many thousands to millions of messages per second).

"In financial markets, microseconds can make the difference. We are pleased to be working with IBM to offer significantly higher throughput and significantly lower latency as proven by our performance tests."

— Patrick Guay,
Senior Vice President of Marketing,
Voltaire

WebSphere MQ Low Latency Messaging, Version 2.0 extends the WebSphere MQ product line with a new transport designed for low-latency, high-throughput messaging. It adds to the existing range of transports a messaging product optimized for the very high-volume, low-latency requirements typical of financial services firms and other industries where speed of data delivery is paramount. Although WebSphere MQ continues to provide the premier solution for rock-solid, assured, time-independent message delivery, the addition of WebSphere MQ Low Latency Messaging to the WebSphere messaging portfolio augments this comprehensive suite of transport protocols to address an increasingly broad range of quality of service (QoS) requirements.

IBM WebSphere MQ Low Latency Messaging delivers:

- *Very high messaging throughput with low latency*
- *One-to-many multicast messaging*
- *Point-to-point unicast messaging*
- *Support for Support for User Datagram Protocol (UDP) and Transmission Control Protocol (TCP)*
- *Positive or negative message acknowledgement*
- *Stream failover for high availability*
- *Flexible, fine-grained message filtering*
- *Traffic rate and congestion control*
- *Robust monitoring of application and network statistics, including internal and external latency*
- *Support for Linux, Windows and Solaris platforms*

High-speed messaging

WebSphere MQ Low Latency Messaging offers the messaging technology previously available as a component of WebSphere Front Office for Financial Markets as an enhanced, separately available product. Its messaging QoS can be used as a stand-alone element of financial markets solutions or in other industries that have similar needs for reliable, high-speed data delivery. WebSphere MQ Low Latency Messaging is also embedded in WebSphere Front Office.

Based on technology developed at the IBM Haifa Research Lab, WebSphere MQ Low Latency Messaging achieves its breakthrough speed by packetizing data efficiently and exploiting Internet Protocol (IP) multicast infrastructure in a daemonless fashion to eliminate network connections. WebSphere MQ Low Latency Messaging provides a multicast transport for high-speed, one-to-many communications using UDP with receiver feedback. Although typical multicast implementations offer only best-effort, unreliable message delivery, the addition of delivery options for receiver feedback enables business-based tradeoffs of speed over reliability or reliability over speed.

WebSphere MQ Low Latency Messaging offers two transports in addition to reliable multicast. The first alternative is a lightweight, point-to-point UDP transport with either positive- or negative-feedback reliability and traffic control features similar to the multicast offering. With positive acknowledgment, all packets are acknowledged, whereas negative acknowledgment provides feedback only if a packet is lost. The second alternative offers reliable, point-to-point, unicast messaging over TCP/IP, in which reliability and traffic control are primarily handled by the TCP protocol. These alternatives provide the ability to deliver a stream of data across a wide area network (WAN) or through a firewall.

Standard multicast transports do not offer congestion-control or failover capabilities, resulting in competition between receivers and between streams with subsequent delays and data loss. WebSphere MQ Low Latency Messaging adds traffic-control features on top of the multicast layer for both multicast and unicast transports. Both multicast and unicast transports include methods to monitor traffic (including transmission rate, losses and retransmissions, and latency) to notify the application of network-congestion problems and to manage these detected problems by handling slow receivers or regulating the transmission rate.

WebSphere MQ Low Latency Messaging supports highly available message transmission. All transports enable high- availability distribution by implementing a number of stream failover policies that allow seamless migration of message transmission from failed to backup processes. The unicast transport enables connection retry and resend as necessary, and heartbeat signals to verify that the connection is kept alive.

WebSphere MQ Low Latency Messaging also offers a range of message-filtering options, including coarse-grained, topic-based and fine-grained filtering. This flexibility allows the application to control the amount of data that is delivered to each application to make the most efficient use of network bandwidth and processing resources.

InfiniBand is a next-generation interconnect standard that offers high transmission rates and scalability. WebSphere MQ Low Latency Messaging delivers support for InfiniBand over IP to enable higher throughput with even lower latency, reduced latency variability and low central processing unit (CPU) consumption

"The performance and interoperability of WebSphere MQ Low Latency Messaging will give our clients even more open access to Reuters systems at greater speed than ever. For the first time, developers will be able to build new applications using IBM messaging software to access content from [Reuters Market Data System] RMDS and to distribute it into other parts of their organizations. This is another example of our key aim to put customers first, and we look forward to working more closely with IBM as it introduces further enhancements to the suite of IBM WebSphere Software products."

—Peter Moss,
Global Head of Enterprise Solutions,
Reuters

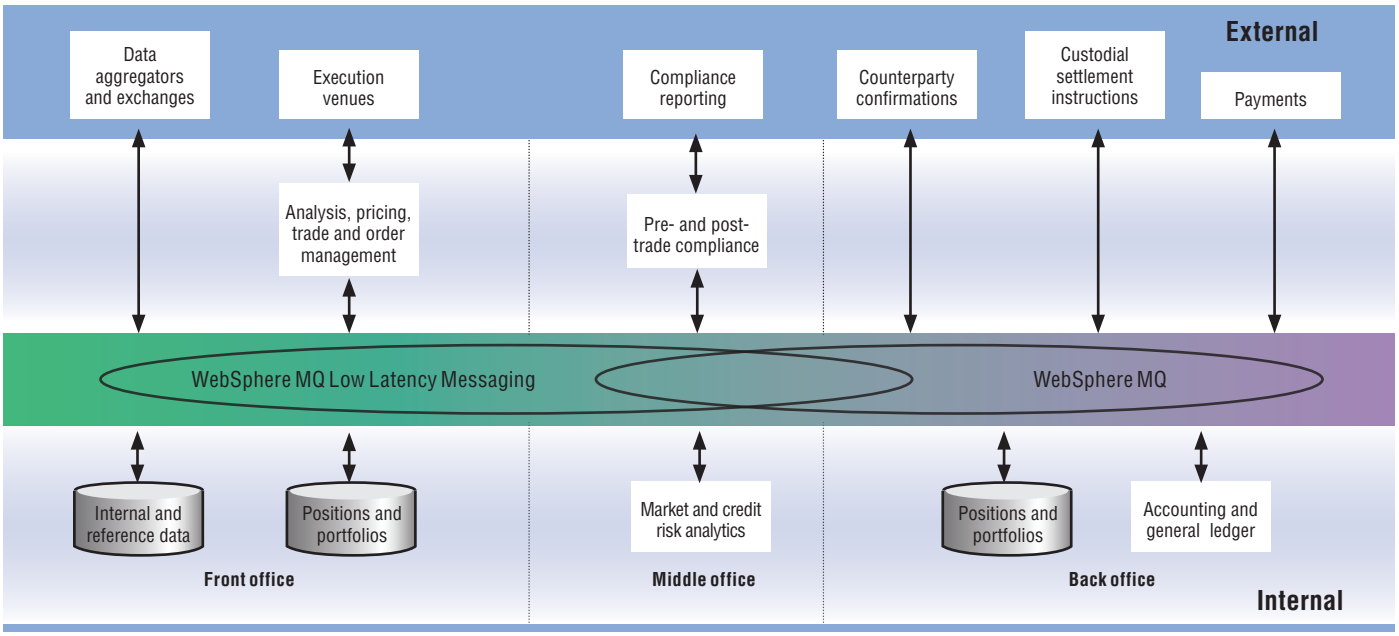

Figure 1. Low-latency messaging and WebSphere MQ in the trade process

Applications of WebSphere MQ Low Latency Messaging

With the rapid growth of algorithmic and model-driven trading and direct market access, financial services firms require superior messaging infrastructures to manage the massive explosion of data to remain competitive. "Firms are turning to electronic trading, in part because a 1-millisecond advantage in trading applications can be worth millions of dollars a year to a major brokerage firm."[1]

As shown in Figure 1, WebSphere MQ Low Latency Messaging can be used almost anywhere within the market-data and trading life cycle. This can include delivery of the following types of information:

- *Market data from exchanges to market-data consumers*
- *Market and reference data within the enterprise to analytic or trading applications*
- *Trade data – such as positions, orders and confirmations – to direct market access and other trading applications*
- *Event notifications for systems monitoring, risk analytics and compliance applications*

WebSphere MQ Low Latency Messaging can help address these common issues:

- Increasing market-data volumes: *Fragmented orders and higher message traffic due to growth in electronic exchanges, high-frequency trading strategies and longer trading hours*
- Pressure on existing messaging solutions: *Inability to handle high volume and throughput requirements*
- Increasing regulatory requirements: *Determination required of best price for all client orders, with execution against that price, for example REG NMS and MiFID*
- Excessive message latency: *Caused when daemon-based messaging solutions insert excessive network connections*
- Dropped messages and network storms: *When high volumes of messages interfere with one another*

WebSphere MQ Low Latency Messaging can extend the advantages already available within WebSphere Front Office for the transport of market data to a broad range of messages for the middle and back offices of financial markets, and to other industries with similar needs. Potential future applications of WebSphere MQ Low Latency Messaging could include transportation, chemical and petroleum processing defense, multimedia and other industries that require high-volume, low-latency, reliable message delivery.

Performance tests conducted in IBM's Haifa Research Lab demonstrate the outstanding performance and latency that can be achieved with WebSphere MQ Low Latency Messaging.[2] The test environment included these components:

- *Two IBM BladeCenter® HS21 servers with two Intel® Xeon® 5130 2 GHz CPUs and 4 GB RAM running Linux® RH4 Update 4 (x86-64, 64-bit instruction set)*
- *Gigabit Ethernet network, including one Cisco switch for IBM BladeCenter H and two internal IBM BladeCenter Gigabit Ethernet expansion cards*
- *Voltaire InfiniBand network, including one Voltaire ISR 9096 switch, Version 4.0.0; two InfiniBand x4 HCA cards (PCI-Express x8); three Voltaire GridStack, Version 4.3.0_4_e2; and four InfiniBand multicast libraries by Voltaire, called MCE, Version 207*

The following results were achieved:

- *About one million 120-byte messages per second on Ethernet*
- *Close to three million 120-byte messages per second on InfiniBand*
- *More than 8 million smaller messages per second, all on common x86 servers on Ethernet*
- *Very low latency of 30 microseconds for 120-byte messages*

Throughput				
Message size [bytes]	Network			
	InfiniBand (MCE)		Ethernet (1 Gbps)	
	msgs/sec	Mbps	msgs/sec	Mbps
12	8250000	774	8360000	784
45	7080000	2490	2500000	888
120	2900000	2730	970000	912
1200	270000	2540	99000	928
12000	26000	2430	9850	923
120000	2940	2750	987	926

Table 1. Throughput[3] Gbps=gigabits per second Mbps=megabits per second

Single-hop latency (microseconds)					
Network	Message size (bytes)	Transmission rate (msgs/sec)			
		10K	100K	500K	1M
InfiniBand (MCE)	45	29	41	51	61
	120	30	40	62	
	1000	36	80		
Ethernet (1 Gbps)	45	49	71	96	147
	120	61	93	171	
	1000	80	253		

Table 2. Average single-hop latency

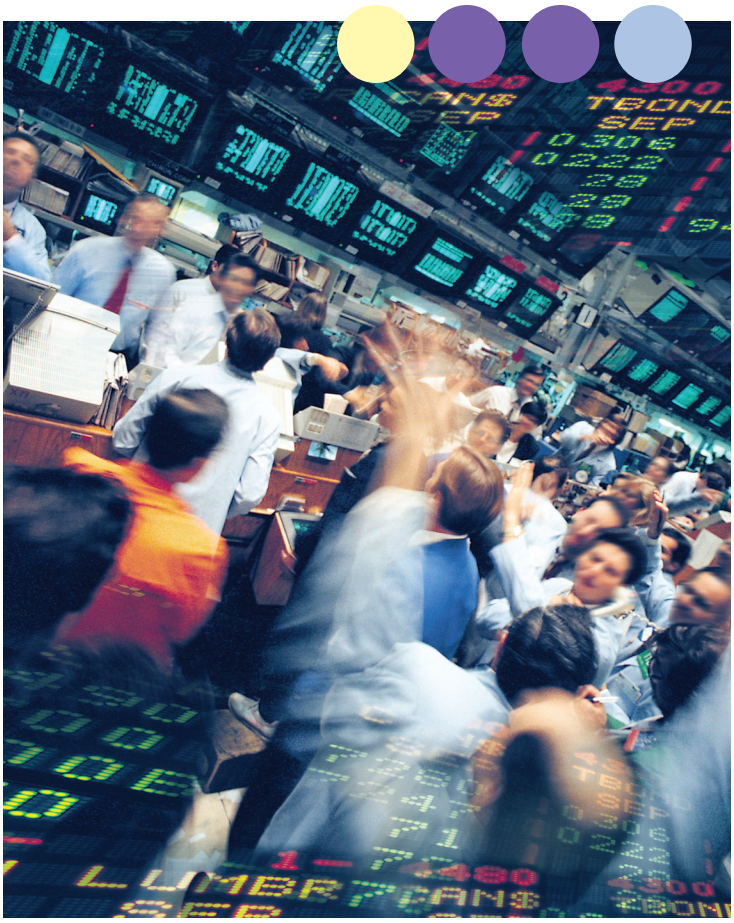

New messaging quality of service

WebSphere MQ Low Latency Messaging is designed to provide the qualities of service that financial services need, as shown in Table 3.

WebSphere MQ family

IBM WebSphere MQ products serve as the backbone for messaging throughout the entire enterprise, both internally and externally. The comprehensive suite of transport protocol and QoS options combined with extensive support and individual computing platforms and runtime environments provide the underpinning for an enterprise service bus (ESB) architecture and speedy adoption of a service oriented architecture (SOA).

For more information

To learn more about how IBM WebSphere MQ Low Latency Messaging, Version 2.0 can help you optimize your investments and reach your business and IT goals, contact your IBM representative or IBM Business Partner, or visit:

ibm.com/software/integration/wmq/llm

	WebSphere MQ	WebSphere MQ Low Latency Messaging
Speed QoS	Time-independent, “as-soon-as-possible” delivery	Ultra-fast delivery in microseconds
Reliability QoS	Ultra-reliable, “bet the business” delivery	Reliable, “know-when-you-lose-data” delivery
Target market	Enterprise messaging	Low-latency messaging
Target industries	All	Financial markets and others with similar QoS needs
Queuing	Yes	No
APIs	JMS, XMS and MQI	Unique, including its own set of APIs optimized for low-latency message delivery
Platform coverage	80 platforms	Linux and Microsoft® Windows® on x86

Table 3. Comparison of IBM messaging transports

IBM WebSphere MQ Low Latency Messaging, Version 2.0 at a glance

Hardware requirements

- Processor: AMD or Intel x86 architecture 32- or 64-bit, or Solaris UltraSPARC
- One processor with a minimum processor speed of 2.0 GHz (dual processor or dual core recommended)
- Minimum memory requirements: 1 GB

Software requirements

One of the following operating system platforms:

- Red Hat Enterprise Linux AS/WS 4 Update 5 (x86 or x86-64)
- Red Hat Enterprise Linux 5 Server or Client (x86 or x86-64)
- SUSE Enterprise Linux 10 SP1 Server or Desktop (x86 or x86-64)
- Microsoft Windows XP SP2 (x86 or x86-64)
- Microsoft Windows Server 2003 SP1+ (x86 or 86-64)
- Microsoft Windows Vista (x86 or x86-64)
- Solaris 10 UltraSPARC (32- or 64-bit)
- Solaris 10 (x86 or x86-64)

Development systems requirements

One of the following compilers:

- Linux libraries compiled with gcc 4.1.1
- Windows libraries compiled with Microsoft Visual Studio 8.0
- Solaris libraries compiled with Sun Studio 11

Performance considerations

WebSphere MQ Low Latency Messaging performance depends on the complexity of the specific environment, volume of data traffic and the data object size.

IBM Corporation
Software Group
Route 100
Somers, NY 10589
U.S.A.

Printed in the U.S.A.
11-07

[1] Martin, Richard. 2007. "Data Latency Playing An Ever Increasing Role In Effective Trading." Attributed to Information Week by Wall Street & Technology. http://www.wallstreetandtech.com/resourcecenters/low-latency/showArticle.jhtml?articleID=199702208 (accessed October 25, 2007).

[2] The performance numbers listed for WebSphere MQ Low Latency Messaging are based on measurements using standard IBM benchmarks in a controlled environment. The actual throughput that any application will experience may vary depending upon considerations such as message size, transmission rate, hardware platform, and network configuration. Therefore, no assurance can be given that an individual application will achieve the throughput or latency stated here. Customers should conduct their own testing. For more detailed performance information, consult your IBM sales representative.

[3] Notes for Table 1:

- When using Ethernet, the throughput of reliable multicast messaging (RMM) is limited by the network speed, which is 1000 Mbps.
- Using a simple load driver shows that the throughput of the MCE library over Infiniband is 2750 Mbps. We can see that, except for the case of 12 bytes messages, the throughput of RMM is almost the same. That is, RMM is limited by the maximum throughput of the MCE library.
- The reported throughput (in the Mbps column) is that of the application. The throughput going out to the network is actually higher due to RMM and network headers, which are appended to each packet.

WSB14003-USEN-00